Estimating Subgraph Frequencies with or without Attributes from Egocentrically Sampled Data

Minas Gjoka[†], Emily Smith[◦], Carter T. Butts[†◇◇]
EECS Dept.[†], Sociology Dept.[◦], Statistics Dept.[◇]
University of California, Irvine
{mgjoka, emilyjs, buttsc}@uci.edu

arXiv:1510.08119v1 [cs.SI] 27 Oct 2015

ABSTRACT

In this paper we show how to efficiently produce unbiased estimates of subgraph frequencies from a probability sample of egocentric networks (i.e., focal nodes, their neighbors, and the induced subgraphs of ties among their neighbors). A key feature of our proposed method that differentiates it from prior methods is the use of egocentric data. Because of this, our method is suitable for estimation in large unknown graphs, is easily parallelizable, handles privacy sensitive network data (e.g. egonets with no neighbor labels), and supports counting of large subgraphs (e.g. maximal clique of size 205 in Section 6) by building on top of existing exact subgraph counting algorithms that may not support sampling. It gracefully handles a variety of sampling designs such as uniform or weighted independence or random walk sampling. Our method can be used for subgraphs that are: (i) undirected or directed; (ii) induced or non-induced; (iii) maximal or non-maximal; and (iv) potentially annotated with attributes. We compare our estimators on a variety of real-world graphs and sampling methods and provide suggestions for their use. Simulation shows that our method outperforms the state-of-the-art approach for relative subgraph frequencies by up to an order of magnitude for the same sample size. Finally, we apply our methodology to a rare sample of Facebook users across the social graph to estimate and interpret the clique size distribution and gender composition of cliques.

1. INTRODUCTION

In a large number of real-world applications it is common to represent systems, structures, or data using graphs *e.g.,* social graphs, web graphs, or protein interaction graphs. In many cases these graphs are difficult to study, most commonly because of their massive size and/or access limitations. As a result, there is a growing body of work [15,22,32] that uses sampling to estimate the properties of such graphs as a step towards understanding them. In this paper, we show how to efficiently produce unbiased estimates of the count of an (optionally maximal) subgraph or induced subgraph of a given form in a graph or digraph from a probability sample of nodes, with or without nodal attribute constraints.

There has been great interest in the research community in counting either statistically over-represented (partial or induced) subgraphs [25], sometimes called network motifs [39], or the full census of induced subgraphs of a given size, called graphlets [42] (typically 3, 4, and 5-node). In the seminal work by Holland and Leinhardt [24], global network structures were shown to be heavily constrained by their subgraph composition; this work also established the enumeration and labeling of graph isomorphism classes

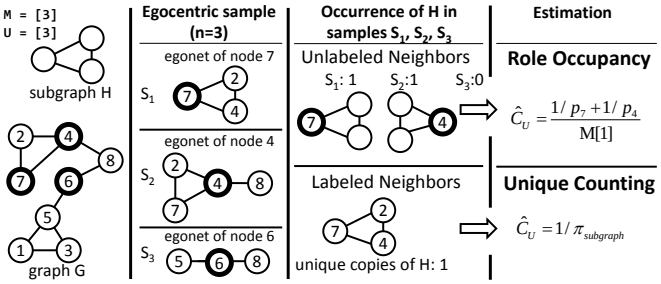

Figure 1: Illustration of the egocentric approach, and the difference between labeled and unlabeled neighbors. The goal is to estimate the subgraph H count in graph G. In this case, egos 7, 4, 6 are sampled from graph G.

to identify specific types of subgraphs, an approach that has since become standard practice in many fields [4,42,54]. [14,26] introduced use of parametric statistical models for networks based on subgraph counts, an approach that is also now in wide use [27]. Parametric network models based on graphlets were introduced by [58], who applied them to social and protein structure networks. Particular classes of subgraphs have also found applications in a number of fields. Cliques, for instance, have been studied in social networks as the foundation for clustering and cohesive subgroups [55]. In addition to social networks [5,28,33], cliques have been analyzed in a diverse range of networks, including those relating to protein structure [19,48], image recognition [47], and written texts [9]. Open or incomplete two-paths have been employed as motifs for the study of brokerage in social and organizational systems [18,45], and k-stars have been used as tools for modeling degree distributions [44]. Vital for all of these applications is the ability to count subgraphs of particular types.

The demand for subgraph counts has spurred the development of a variety of algorithms. They can be classified in two categories: those that use exact counting [11,20,23,29,43] and those that use sampling [6,30,41,53,56]. Exact algorithms require knowledge, processing, and storage of the entire graph. In contrast, we provide estimators for subgraph counts using only samples of network data. More importantly, our methods *extend support for sampling* to all exact subgraph enumeration techniques that do not support it by design. In comparison with existing sampling methods, our methods additionally support annotated attributes, and either induced or non-induced subgraphs. In Sec. 5 we show that our methods outperform the state-of-the-art approach of [53] by up to an order of magnitude for the same sample size.

Our estimation techniques for subgraph counts employ an

egocentric approach [55], illustrated in Fig. 1. We first collect a probability sample of nodes ("egos") from the target graph. In our example we sample nodes 7, 4, 6. Then, we collect the egonet of each sampled node, which consists of the neighbors of the node and the edges between these neighbors. Next, we use an existing enumeration tool to calculate the exact subgraph count in each sampled egonet. Our approach supports the absence or presence of unique neighbor labels in the egonets. Fig. 1 shows the implications of either possibility in the partial count. Last, depending on the existence of neighbor labeling we apply the Role Occupancy or Unique Counting estimation method to combine the subgraph counts in individual egonets and estimate the subgraph count for the whole graph in an unbiased manner.

Our approach has several benefits. The first obvious benefit is that it allows us to estimate the subgraph count for a given isomorphism class from an unknown graph, as long as we have a sampling primitive that reveals the neighbors of a selected node in that graph. The second benefit is that it can be used to estimate the subgraph count in a fully known massive graph. Our approach decomposes a large problem into many smaller problems that can be independently computed, hence making estimation embarrassingly parallel. Another benefit is the ability to estimate the subgraph count in the absence of unique neighbor node labels *e.g.,* due to privacy-sensitive network data or data collection limitations. Finally, our techniques can be employed with data collected using standard techniques in both online (*e.g.,* random walk or user ID sampling [16]) and offline (*e.g.,* survey instruments [37]) settings.

In summary, we make the following contributions:

- We present two unbiased estimation methods to efficiently estimate the frequency of a subgraph from an egocentric probability sample of nodes. Our methods support subgraphs of any order so long as they are contained within an egonet, including variations such as undirected/directed, induced/non-induced, maximal/general cases. Our methods also allow for counting of subgraphs that are differentiated by nodal attributes. Our first method, Role Occupancy, is applicable for egocentric data whether or not the neighbors of sampled nodes can be uniquely identified across draws (*unlabeled*), and our second method, Unique Counting, provides additional statistical power where neighbors are uniquely labeled.

- We evaluate our methods on a variety of real-world graphs and sampling methods and provide suggestions for their use. Unique Counting is shown to have on average smaller error than the Role Occupancy method. However that comes with additional space complexity, which can be quite significant depending on the subgraph of interest. We show that choice of sampling method affects the estimation error, with decorrelated random walks and weighted independence sampling having the smallest error.

- We apply our methodology to a sample of Facebook (FB) users to estimate the clique size distribution and gender composition of cliques across the social graph. Through our analysis we discover evidence for strong heterogeneity in the makeup of cliques.

The structure of the remainder of the paper is as follows. Section 2 reviews related work. Section 3 presents the definitions and basic concepts. Section 4 presents our estimation methodology. Section 5 presents simulation results on real-life fully known graphs. Section 6 applies our estimators to samples collected from Facebook. Finally, Section 7 concludes the paper.

2. RELATED WORK

Egocentric sampling is a widely used method for gathering network data [10,36]. This method samples individual nodes and then expands to include their neighborhoods. While this procedure does not necessarily describe the structure of the entire network, it can yield representative samples of the network [36]. Standard random sampling methods can be used to obtain egocentric network data and generalize the results to a larger population [36]. Examples of applications of egocentric sampling procedures include the network items in the General Social Survey [8] and networks obtained through crawling online social networks [16].

Numerous algorithms have been developed to calculate subgraph counts. They can be classified in two categories, those that use exact counting [11,20,23,29,43] and those that use sampling [6,30,41,53,56]. Exact algorithms require knowledge, processing, and storage of the entire graph. In contrast, we provide estimators for subgraph counts using only samples of network data. It is important to note that our methods *extend support for sampling* to all previous exact subgraph enumeration techniques that do not support it by design. This is possible by applying exact counting in each egonet and combining the egonet calculations with our Role Occupancy or Unique counting estimators, as described below. Because of this property, the methods introduced here can be considered complementary to (rather than competitive with) exact counting methods.

mFinder [30] was the first algorithm for the estimation of subgraph count using edge sampling. However, it is computationally intensive and scales poorly with the size of subgraphs (up to 6-node). FANMOD [56] uses a node sampling approach and has improved computational complexity vs mFinder (up to 8 nodes). Compared to our approach, mFinder and FANMOD are limited by the fact that they require a uniform independence sample of edges and nodes (respectively) for unbiased estimation. MODA [41] is another sampling algorithm that uses a pattern growth tree approach. While MODA is not as fast as FANMOD, it is able to find larger motifs [57]. However, the method does not provide guarantees about bias. GUISE [6] is an algorithm that uses MCMC to sample graphlets, estimating 3-, 4-, and 5-node connected induced subgraphs. This can be used for constructing a graphlet frequency distribution and is much faster than a counting-based approach. Finally, [53] developed two algorithms, PSRW and MSS, to estimate induced subgraph counts. PSRW uses appropriately re-weighted random walk samples and is shown in simulations to have lower estimation errors than FANMOD and GUISE, the latter due to not rejecting samples. MSS, a generalization of GUISE, jointly estimates induced subgraphs of size $k-1$, k, and $k+1$ for $k \geq 4$. Unlike GUISE, PSRW, and MSS, our methods additionally support annotated attributes, and either induced or non-induced (and maximal or non-maximal) subgraphs. In Section 5 we show that for the estimation of 3-node directed, and 4-node undirected subgraphs our methods outperform PSRW [53] by up to an order of magnitude

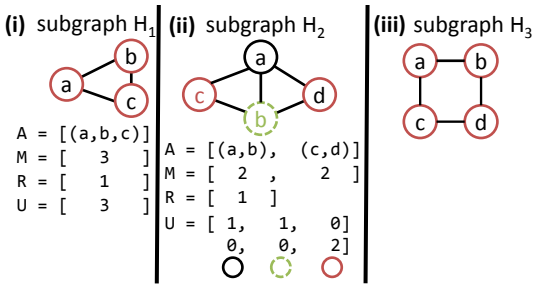

Figure 2: Calculation of A, M, R, U for supported subgraphs H_1, H_2. An example of a non-supported subgraph is H_3.

for the same sample size.

In our recent work [17], we presented statistically principled estimators that count clique subgraphs of all sizes in a graph using ego-centric sampled network data. In this paper, we extend those estimators to support counting all types of subgraphs contained within an egonet. We also examine the effect of network sampling method on estimation error, compare against the state-of-the-art technique for subgraph counting, and demonstrate an application of our estimators to a rare sample of Facebook egonets collected across the whole social graph.

3. DEFINITIONS AND BASIC CONCEPTS

Let $G = (V, E)$ be a undirected or directed graph, with $N = |V|$ nodes and $|E|$ edges. We associate with the vertices of G a vector X of discrete states, features, or attributes (referred to generically as covariates), such that X_i is the state of the i-th vertex of G. Without loss of generality, we denote the possible states of X_i by the integers $1, 2, ..., p$; in the case for which $p \to \infty$, it will be possible for our purposes to limit ourselves to the subset of observed states (of which there can be at most N). In a typical social network context, X will represent a categorical or ordinal covariate such as gender or level of education, or else a Cartesian product of such covariates (e.g., gender by level of education, gender by race, etc.). It is even permissible for X to indicate positional characteristics (e.g., having concurrent partnerships), so long as these are available via the measurement scheme described below. In some cases, we may be interested in graphs whose vertices lack distinguishing characteristics. In such instances, we take X to be a vector of 1s. We refer to the attributes X of V as *annotations*, and say that a graph is *annotated* when it is associated with some attribute vector X; a graph without such an association is said to be *unannotated*.

3.1 Objective

Our goal is to count the number C_U of subgraphs or induced subgraphs within G of a given attribute composition U that are isomorphic to a particular graph; we define this formally as follows. Let H be an unannotated graph of order h, representing the structure to be counted, with the directedness of H matching that of G. Denote by A the automorphism orbits of H, with M the multiplicities of those orbits (such that M_i is the number of positions within orbit A_i of H). For measurement purposes, we will be particularly concerned with orbits whose members are adjacent to the rest of H, for which we introduce specific terminology. For undirected H, we refer to an orbit A_i as being a *spanning orbit* if any vertex $v \in A_i$ is adjacent[1] to all $v' \in H \setminus v$. In the directed case, we likewise refer to A_i as being a *semi-spanning orbit* if all v in A_i have semi-edges to all $v' \in H \setminus v$; an *out-spanning orbit* if all v in A_i are adjacent to all $v' \in H \setminus v$; and an *in-spanning orbit* if all $v' \in H \setminus v$ are adjacent to all $v \in A_i$. Obviously, a graph H containing a spanning or semi-spanning orbit has semi-diameter less than or equal to 2 (although its diameter in the directed case may be larger), though it may contain any number of vertices. It is this family of graphs with which we are concerned.

While H itself determines the structural form of the subgraph to be counted, we are also concerned with the attributes (or composition) of the vertices within it. Let a be the number of automorphism orbits of H. We then define a *composition matrix*, U, for H to be a $a \times p$ matrix of non-negative integers whose row sums are equal to M (and, therefore, whose total sum is equal to h), with U_{ij} indicating the number of members of orbit A_i whose vertex states are equal to j. Fig. 2(ii) shows an example for subgraph H_2 that consists of three node attributes represented by colors red, green and black. We calculate A, M, and U as follows. H_2 has two automorphism orbits : the first consists of nodes a, b whereas the second consists of nodes c,d. Hence the multiplicity of both automorphism orbits is two. In the composition matrix U we assign one column to each attribute. We then count the number of nodes of a given attribute in each automorphism orbit.

For the unannotated case, we have $p = 1$ and $U = M$ (as an $a \times 1$ matrix); likewise, for the case in which $a = 1$ (e.g., cliques), U is a single row-vector corresponding to the number of H-members having each X state. In the trivial case of unannotated H with $a = 1$, U is simply equal to h, the order of H, as seen in the example of Fig. 2(i). Typically, however, U will be more elaborate, and it provides the general mechanism by which we will indicate the specific instances of H whose prevalence we seek to estimate.

3.2 Measurement Assumptions

We presume that our data comes in the form of an *egocentric network sample* $Y'_1, ..., Y'_{n'}$, which is a probability sample of n' egonets from G; since the data may be sampled with or without replacement, we denote the unique elements of the sample in arbitrary order by $Y_1, \ldots, Y_n$, with $n' \geq n$. We define the complete egocentric network (or egonet) of vertex v_i (ego) to be v_i together with an appropriately chosen neighborhood of v_i in G (alters), as well as all edges among the indicated vertices (both ego and alters). The neighborhood over which the egonet is defined (denoted $N^e(v_i)$) is a condition of the measurement process, and may vary; we consider four scenarios:

- G is undirected, and $N^e(v_i)$ includes all $v' \in V$ such that v_i and v' are adjacent (*i.e.,* $N^e(v_i) = Neigh(v_i)$).
- G is directed, and $N^e(v_i)$ includes all $v' \in V$ such that v_i, v' is not a null dyad (*i.e.,* $N^e(v_i) = Neigh^+(v_i) \cup Neigh^-(v_i)$).

[1]Note that H may contain loops, and we do not require that a vertex be tied to itself to belong to a spanning orbit.

- G is directed, and $N^e(v_i)$ includes all $v' \in V$ such that v_i is adjacent to v' (*i.e.,* $N^e(v_i) = Neigh^+(v_i)$).
- G is directed, and $N^e(v_i)$ includes all $v' \in V$ such that v' is adjacent to v_i (*i.e.,* $N^e(v_i) = Neigh^-(v_i)$).

In addition to the egonet structure, we assume that the relevant vertex attributes (i.e., X values) are also known for both egos and alters. The egonet of v_i, then, can be represented formally by the vertex set, edge set, value tuple

$$Y_i = (V_i, G[V_i], X[V_i]), \tag{1}$$

where $V_i = N^e(v_i) \cup v_i$, $G[V_i]$ is the subgraph of G induced by V_i, and $X[V_i]$ is the subvector of X corresponding to the vertices of V_i. Our sample is represented by a vector S of length n whose elements index the unique sampled egos, and a vector p of length n whose elements index the *inclusion probability* for each sampled ego. Thus, if $S_i = j$, this implies that the ith member of the sample is v_j, and v_j was included in the sample with probability p_i. We use this notation in the development that follows.

The choice of egonet neighborhood, N^e, determines the subgraphs or induced subgraphs that may be counted using our approach. In the case of undirected graphs, the subgraphs that may be counted are all H containing at least one spanning orbit. For directed graphs, three possibilities exist: with $N^e(v_i) = Neigh^+(v_i) \cup Neigh^-(v_i)$, we may count any H containing at least one semi-spanning orbit; if $N^e(v_i) = Neigh^+(v_i)$, we may count any H containing at least one out-spanning orbit; and, finally, if $N^e(v_i) = Neigh^-(v_i)$, we may count any H containing at least one in-spanning orbit. Intuitively, this is because we require at least one vertex within each copy of H that will contain that copy within his or her egonet - otherwise, we cannot measure it. We note that additional constraints on Y_i beyond those discussed here (e.g., degree constraints) may place additional constraints on measurable H; here, we consider the case in which sampled egonets are measured completely and exactly.

3.3 Sampling methods

As noted above, we assume that our egonets comprise a probability sample of the egonets in G; that is, (i) we can treat each ego as being included in the sample with known probability, and (ii) the probability of sampling any vertex v is positive for all $v \in V$. Our estimation supports many sampling methods, including the following.

Uniform Independence Sampling (UIS), where nodes are sampled independently with equal probabilities.

Weighted Independence Sampling (WIS), where nodes are sampled independently with probability proportional to a known weight $w(v)$.

Simple Random Walk (RW) [35] selects the next-hop node v uniformly at random among the neighbors of the current node u. On a connected and aperiodic graph, RW samples node v with a limiting distribution proportional to its degree $\deg(v)$.

Sampling may occur with or without replacement; we indicate these distinctions where they affect estimation. Note also that, in practice, samples drawn using link-trace methods e.g. Metropolis Hastings Random Walk [16] or Weighted Random Walk [34], may closely approximate UIS or WIS, and may be employed as well. We provide an example of this approach in Section 6.

3.4 Neighbor Labeling

When egonet S_i of ego i is sampled, it may or may not be possible to uniquely identify i's neighbors (in the sense of knowing, e.g., whether $v \in Y_i$ also belongs to some Y_j). When such identification is possible, we say that the sample is *labeled*, otherwise denoting it as *unlabeled*. Fig. 1 shows the effect of labeling in an example graph in which egos 7, 4, and 6 are sampled from graph G. If the sample is labeled, we can discern that the three sampled egonets contain only 1 unique copy of subgraph H: $\{7, 2, 4\}$. In the unlabeled case, however, we know only that egonet 7 contains one instance of subgraph H, and egonet 4 contains a second instance. As we will show, estimation is possible in both cases; however, labeled samples provide additional information that can be leveraged to reduce sampling error.

4. ESTIMATION

Given a choice of H and U, we estimate the number of subgraphs or induced subgraphs having the appropriate structure within G,X from an egocentric network sample $Y'_1, .., Y'_{n'}$ with unique elements $Y_1, \ldots, Y_n$; in the case of sampling with replacement, n may be less than n'. We propose two families of approaches: one, based on role occupancies, that does not require unique identification of neighbors; and one, based on unique subgraph counts, that leverages them.

4.1 Role Occupancy Method

Given H, A, and the choice of egocentric neighborhood used in the measurement of the graph, we define a given orbit of A to be *observable* iff membership of some v_i in such an orbit implies that the associated subgraph (isomorphic to H) must be contained in v_i's egonet. As noted above, our method is applicable to any H with at least one such orbit (given choice of N^e). Let R be a vector containing the indices of the observable orbits in A, such that A_{R_i} is observable for all i; denote the length of this vector by r (the number of observable roles). For example in the subgraph H_2 of Fig. 2(ii) only orbit 1 that contains nodes a, b is an observable orbit. We then define the (observable) *role occupancy degrees* of a vertex, v_i, to be the vector d_{iU} such that d_{iU_j} is the number of U-composition subgraphs or induced subgraphs of type H for which v_i occupies role R_j. It is important to note that, given Y_i, d_{iU_j} can be determined exactly for all observable roles, as this is the basis of our technique.

Given the above, we may likewise define the role occupancy degree sums,

$$D_{U_j} = \sum_{i=1}^{N} d_{iU_j}, \tag{2}$$

which aggregate role occupancies across the graph. Since each copy of H with composition U has exactly M_{R_i} nodes in the ith observable role, it follows that D is deterministically related to the count C_U of such graphs. Let $m_i = M_{R_i}$ be the multiplicity of the ith observable orbit. Then,

$$\sum_{j=1}^{r} D_{U_j} = C_U \sum_{j=1}^{r} m_j, \tag{3}$$

and so

$$C_U = (\sum_{j=1}^{r} D_{U_j})/(\sum_{j=1}^{r} m_j). \quad (4)$$

Since the sum of multiplicities is a constant, we can obtain an unbiased estimate of C_U immediately from any unbiased estimators of the corresponding degree sums.[2] A natural way to achieve this is via Horvitz-Thompson (H-T) estimation. Define

$$\widehat{D_U} = \sum_{i=1}^{n}\sum_{j=1}^{r} d_{S_i U_j}/p_i \quad (5)$$

to be the H-T estimator of $\sum_{j=1}^{r} D_{U_j}$, where p_i is the node inclusion probability of i (i.e., the probability of i appearing in the sample at least once). It then follows that

$$\widehat{C_U} = \widehat{D_U}/(\sum_{j=1}^{r} m_j) \quad (6)$$

is an unbiased estimator of C_U. Variance estimates can be worked out from H-T theory in the usual manner, exploiting the fact that $\mathrm{Var}(\widehat{C_U}) = \mathrm{Var}(\widehat{D_U}/(\sum_{j=1}^{r} m_j)^2)$.

For designs where the probability that any two observed nodes, j and k, are both included in the sample is known, unbiased estimators of the subgraph estimator variance are given by the general form [3]

$$\widehat{Var}(\widehat{C_U}) = \sum_{j=1}^{n}\left(\frac{1}{p_j^2} - \frac{1}{p_j}\right)\left(\frac{\sum_{i=1}^{r} d_{S_j U_i}}{\sum_{i=1}^{r} m_i}\right)^2 + 2\sum_{j=1}^{n}\sum_{k=j+1}^{n}\left(\frac{1}{p_j p_k} - \frac{1}{p_{jk}}\right)\left(\frac{\left[\sum_{i=1}^{r} d_{S_j U_i}\right]\left[\sum_{i=1}^{r} d_{S_k U_i}\right]}{(\sum_{i=1}^{r} m_i)^2}\right), \quad (7)$$

p_{jk} above is the probability of both j and k appearing in the sample. For designs such that p_{jk} cannot be readily determined, the generalized H-T estimators of form *Eq.*(8) below can be employed.

An important special case arises when sample inclusion probabilities are unequal and known only up to a constant factor (i.e., some $w_j \propto p_j$). Given that joint inclusion probabilities are not available here, an adaptation of the Brewer and Hanif (B-H) variance estimator [7] leads to the following general form:

$$\widehat{Var}(\widehat{C_U}) = \left(\frac{N-n}{n(n-1)N}\right)\sum_{j=1}^{n}\left(\frac{n\sum_{i=1}^{r} d_{S_j U_i}/w_j}{\sum_{k=1}^{n}\sum_{i=1}^{r} m_i/w_k} - \widehat{C_U}\right)^2 \quad (8)$$

The B-H estimator is generally biased upward [49], and is hence a conservative estimate of measurement error, but does not require joint inclusion weights.

4.2 Unique Counting Method

When alters can be uniquely identified across egonets, it becomes possible to estimate C_U by counting unique copies of H. Let c_u be the count of unique copies of H with composition U in the sample, such that an ego belongs to an observable role in each copy. An H-T estimator of C_U is then immediately given by

$$\widehat{C_U} = \sum_{i=1}^{c_u} 1/\pi_i \quad (9)$$

where π_i is the probability that the ith unique copy appears in the sample.

As an H-T estimator, the counting estimator is unbiased, and it will generally be more efficient than the role occupancy estimator. It may, however, be much harder to implement, and in particular its space complexity is much greater. Obviously, it also requires labeling information. Given these tradeoffs, both estimators are of potential merit in practice.

4.3 Inclusion Probabilities

We have provided estimators of subgraphs counts for either labeled or unlabeled egonet samples. To use them, it remains only to determine the inclusion probabilities of nodes or subgraphs (p and π, above). These quantities depend on the sampling design; we here provide examples for some common and important cases for sampling of OSNs in particular but also other arbitrary graphs.

Node inclusion probabilities.

The simplest case for node inclusion probabilities is that in which egos are sampled uniformly at random from the population (UIS). The inclusion probabilities depend upon the total number of samples drawn ($n' \geq n$), and whether samples are drawn with or without replacement. In the with-replacement case, an arbitrary node j fails to be selected on any given draw with probability $1 - 1/N$, and hence is ultimately included with total probability $p_j = 1 - (1 - 1/N)^{n'}$. When sampling is performed without replacement, a total of $n' = n$ of the N available nodes are drawn, any of which could be j. The resulting inclusion probability is thus simply $p_j = n'/N$.

When the probability of inclusion on any given draw is unequal, total inclusion probabilities may depend on the details of the sampling mechanism. In the common case of independent with-replacement sampling with unequal probabilities (WIS), the probability of including node j can be determined from the probability of obtaining j on any *given* draw, p'_j, by $p_j = 1-(1-p'_j)^{n'}$. Without-replacement inclusion probabilities with unequal are not easily summarized, but computational tools such as [50] can be employed to obtain them.

In some cases the per-draw inclusion probability may be unequal and known only up to a constant factor (i.e., $w'_j \propto p'_j$). This situation is common in e.g. random walk sampling of OSNs, where vertices are often sampled (approximately) independently with replacement, proportional to degree. In such cases, approximating p'_j by the Hansen-Hurwitz [21] estimator $p'_j \approx w'_j \left(\sum_{k=1}^{n'} 1/w'_k\right)/(n'N)$ (where the sum is over all observations, including repetitions) is a practical alternative.

Subgraph inclusion probabilities.

The probability of sampling a subgraph is equal to the probability that at least one member of at least one observable role in the copy is selected as an ego. The total number of such opportunities is $\sum_{j=1}^{r} m_j$; the resulting inclusion,

[2] Note that sums of unbiased estimators are unbiased estimators of the corresponding sum, a property that follows from linearity of expectation.

[3] This is a direct application of Eq. 6 of [49], p54.

however, may depend on which egos occupy the roles in question. In the case of uniform sampling with replacement (with total draws n'), the inclusion probability is

$$\pi_i = 1 - (1 - (\sum_{j=1}^{r} m_j)/N)^{n'}, \quad (10)$$

while the corresponding without-replacement probability is

$$\pi_i = 1 - \prod_{k=0}^{n'} (N - k - \sum_{j=1}^{r} m_j)/(N - k). \quad (11)$$

When nodes are drawn non-uniformly, the situation can be more complex, but one case is fairly straightforward. Let p'_j be the per-draw sampling probability for the jth member of copy i, under non-uniform independence sampling of nodes. Then the total inclusion probability for the ith copy of H is

$$\pi_i = 1 - (1 - \sum_{j=1}^{\sum_{k=1}^{r} m_k} p'_j)^{n'}. \quad (12)$$

4.4 Implementation considerations

The estimation of the absolute subgraph count requires the graph size N and the enumeration of subgraphs for each sampled egonet. In the cases when N is not known a priori, [31] provides estimators that work with sampled network data. In general, the exact enumeration of subgraphs is a hard problem. In our approach we avoid enumeration over the whole graph, enumerating only within each sampled egonet. It is important to note that (1) each separate computation can be accelerated because our estimators require subgraph counts only for each ego net; this changes the typical complexity of subgraph counting from $O(f(N))$ to $O(f(\deg_m))$ where $\deg_m$ is the maximum degree (often constant in N). Likewise, (2) the computation on each egonet is independent and thus can be parallelized. Additionally, we can always use the fastest state-of-the-art enumeration tool for the given subgraph type, which speeds up the estimation process. Further, our approach can build on top of existing motif enumeration tools that do not support sampling. For example, in the simulation section we use the maximal clique listing method by [12], the 4-node and 5-node graphlet count method Orca by [23], and a customized subgraph search count by Sage [46] for the 3-node directed subgraphs. Neither of these tools supported subgraph counting with sampling and all of them employ the fastest known algorithms to enumerate the subgraphs that they support.

The space complexity of the Role Occupancy method is $O(n)$, where n is the number of unique egonets sampled. On the other hand, the Unique Counting method requires $O(C_U)$ which can be quite large depending on the subgraph type and the graph structure. To give an example, there are ∼627 trillion distinct 5-node "star" subgraphs in the network "web-Google" and storing them is impractical in this case.

5. PERFORMANCE EVALUATION VIA SIMULATED SAMPLING

In this section, we evaluate the performance of the estimators with labeled and unlabeled neighborhoods (i.e., unique counting versus role occupancy) via simulated sampling from real-world datasets. Our results shed light on the relative advantages of these estimators for counting the frequency of a subgraph using sampling.

Dataset	$\|V\|$	$\|E\|$ undirect.	|E| directed	Max. Degree	Maxim. Cliq. Sz.
ca-CondMat [3]	21 362	91 282	-	279	26
FB:New Orl. [52]	63 392	816 884	-	1 098	30
soc-Slashdot [3]	77 360	469 179	828 161	2 539	26
soc-sign-Epin [3]	119 129	704 265	833 390	3558	94
email-EuAll [3]	224 832	340 795	394 400	7 636	16
amazon0601 [3]	403 364	2 443 311	3 387 224	2 752	11
web-Google [3]	855 802	4 291 350	5 066 842	6 332	44
roadnet [3]	1 087 561	1 541 512	-	9	4
youtube-links [40]	1 134 890	2 987 623	4 942 035	28 754	17
FB:UCSD [51]	14 948	443 215	-	2 165	43
FB:UVA [51]	17 196	789 308	-	3 182	42

Table 1: Empirical topologies used in Sec. 5

5.1 Datasets

Table 1 lists the empirical networks that we use in our evaluation study. It includes several online social networks, an email communication graph, a co-authorship network, a transportation topology and a web graph. Some of the networks have directed edges whereas others have only undirected edges. The former networks are treated as both directed and undirected, depending on the directedness of the subgraph that we are interested in counting C_U. For each network we keep the largest connected component[4] and we list the number of nodes, number of undirected and directed edges, and the maximum degree and maximum clique size in the undirected graph. The numbers of nodes in the graphs range from tens of thousands to millions.

The list of networks consists of two groups. The first group contains no attributes whereas the second group contains several node attributes. We have selected the node attribute "gender" to estimate C_U.

5.2 Error Metrics

We measure the error of an estimator $\widehat{C_U}$ with respect to its real value C_U over k simulation iterations using the Normalized Mean Absolute Error (NRMSE) as follows: $NRMSE(\widehat{C_U}, C_U) = \frac{\sqrt{1/k \cdot \sum_{i=1}^{k} (\widehat{C_U} - C_U)^2}}{C_U}$. In some cases we may want to estimate the count of more than one subgraph e.g. all measurable 5-node undirected subgraphs (see Fig. 6c). We summarize the error over all estimations using the Normalized Mean Absolute Error (NMAE), defined as: $NMAE(\widehat{\vec{x}}, \vec{x}) = \frac{\sum(|\widehat{x}_i - x_i|)}{\sum |x_i|}$ where $\vec{x}$ and $\widehat{\vec{x}}$ are the vectors that correspond to the real and estimated values. $NMAE$ returns the absolute estimation error relative to the true value, averaged over every point in the vector.

5.3 Results

5.3.1 Role Occupancy (RO) vs Unique Counting (UC)

We will first compare the performance of our proposed sampling methods, Role Occupancy and Unique Counting. The subgraphs of interest in this comparison will be all order-i maximal cliques.[5]

[4]When the graph is treated as directed we keep the largest weakly connected component.

[5]A clique is a complete induced subgraph. A clique that contains i vertices is called an order-i clique. A clique is

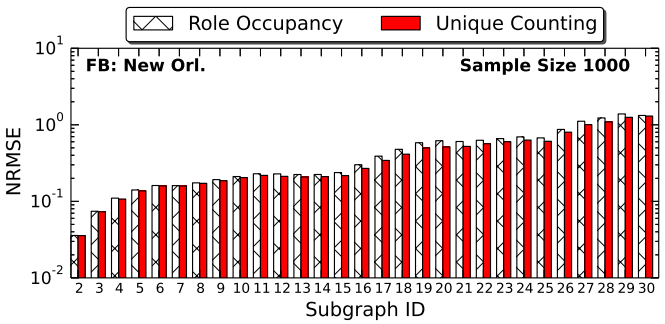

Figure 3: (FB: New Orleans) Subgraph ID i corresponds to the maximal clique of size i. NRMSE calculated over 1000 simulated iterations. 1000 egonets are sampled uniformly without replacement.

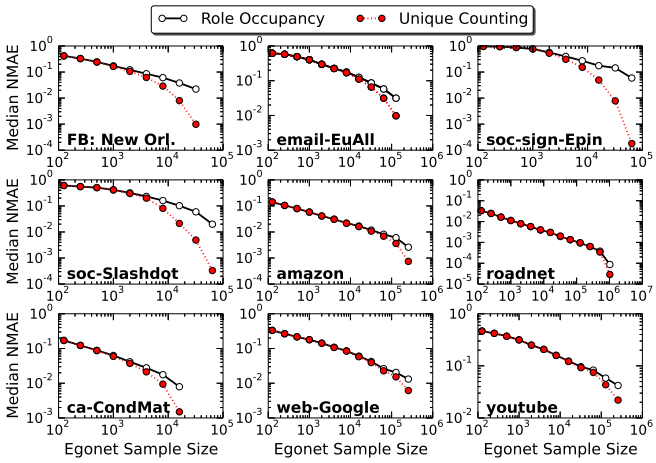

Figure 4: Clique order distribution for real-world topologies. Median NMAE for the estimation of all order-i maximal cliques calculated over 1000 simulated iterations, as a function of the sample size n (uniform sampling without replacement).

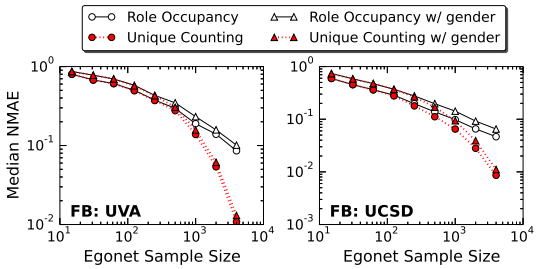

Figure 5: Clique order and composition distributions for real-world topologies. Median NMAE for the estimation of all order-i maximal cliques distinguished by gender calculated over 1000 simulated iterations, as a function of n (uniform sampling without replacement).

Fig. 3 shows the NRMSE for the estimation of all cliques for the topology "FB: New Orleans" (from size 2 up to the maximum clique size 30) using a uniform egocentric sample of size $n = 1000$. The first observation is that NRMSE is higher for larger clique sizes, probably because those are encountered less often. The second observation is that, as expected, Unique Counting is slightly better than Role Occupancy although the difference is very small for this sample size.

To get a better understanding of comparative performance, in Fig. 4 we plot the median NMAE of the RO and UC estimators for all order-i maximal cliques on various real-world topologies as a function of sample size. We vary n from 125 to the total size of each graph, allowing us to observe the effects of saturation on measurement error. We note that for smaller sample sizes, the RO and UC estimators perform equally well. Beyond a threshold sample size, however, the UC begins to substantially outperform the RO estimator (reflecting the additional information associated with vertex labels). We use Table 2 to better interpret these results and shed some light on the causes of the "threshold" behavior. Table 2 contains for each topology and egonet sample size the average % of all nodes and % of all edges when both egos and neighbors are included. We observe that the UC "breakaway" threshold varies for different graphs even when taking into account the total % of nodes and edges . As an example, the threshold for the network "soc-sign-Epinions" is at $n \approx 4\,000$, corresponding to $\approx 18.1\%$ of all nodes being and 34.8% of all edges being contained in some egocentric network sample on average (over 1000 simulations). On the other hand, the threshold for the network "amazon0601" is at $n = 64\,000$, at which point 80.9% of all nodes and 63.3% of all edges have been captured by some egocentric sample on overage. While saturation aids the UC estimator relative to the RO estimator, the degree of saturation required varies markedly.

Next, we examine how the addition of node attributes affects the estimation of counts for subgraphs annotated with attributes. For that reason, we simulate the estimation of all order-i maximal cliques that are distinguished by the "gender" attribute [6] in Facebook networks UCSD and UVA (see Fig. 5). Due to the size and density of these topologies, the egonet sample size is set between $15 - 4\,000$. Table 2 shows the values for the mean % of nodes and % of edges for these egonet sample sizes. As expected from the larger number of values (and smaller counts), estimation of the clique composition distinguished by gender is at least as hard as the estimation without gender. Depending on the composition of the attributes, the estimation w/gender ranges from being indistinguishable (see FB:UVA) or slightly worse than w/out gender (see FB:UCSD).

Our results show clear returns to the use of labeled neighborhoods where possible: the UC estimators perform on average as well or better than the RO estimators in all cases. However, to count the distinct subgraphs the UC estimator needs additional space as discussed in Section 4.4. Depending on the topology and the subgraph of interest, the amount of space required to implement the UC estimator might be considerably high. For example, the estimation of all order-i maximal cliques with labeling for the Facebook '09 data samples in Section 6 requires space that is at least in the order of hundreds of GBytes. Additionally, in some settings (e.g., due to privacy or data collection limitations, particularly offline) is not possible to obtain information on neighbors' identities. In all these cases, our simulations suggest that the RO estimator can still provide excellent per-

said to be *maximal* if no higher order clique contains it.

[6]For example, we estimate all 4 types of order-3 cliques: with 3, 2, 1, or 0 males and corresponding females.

Dataset	% sampled	Egonet Sample Size						
		250	**500**	**1K**	**2K**	**4K**	**8K**	**16K**
FB:New Orl.	% nodes	9.11	16.16	26.61	40.04	55.25	70.56	84.53
	% edges	5.53	10.55	19.08	32.12	49.33	67.82	84.27
email-EuAll	% nodes	0.38	0.76	1.37	2.53	4.65	8.75	16.35
	% edges	0.47	0.97	1.86	3.58	6.70	12.25	21.61
soc-sign-epinions	% nodes	2.23	4.05	6.86	11.31	18.14	28.61	43.77
	% edges	4.31	8.36	14.21	23.30	34.84	49.10	64.32
soc-slashdot	% nodes	3.60	6.51	11.15	18.31	28.71	43.14	61.75
	% edges	1.68	3.31	6.20	11.38	19.90	33.39	53.15
amazon	% nodes	0.80	1.60	3.15	6.15	11.72	21.59	37.48
	% edges	0.42	0.85	1.69	3.35	6.56	12.60	23.37
web-google	% nodes	0.32	0.62	1.23	2.39	4.61	8.69	15.85
	% edges	0.33	0.65	1.31	2.57	5.02	9.61	17.79
youtube links	% nodes	0.13	0.26	0.52	0.98	1.88	3.51	6.46
	% edges	0.10	0.21	0.46	0.85	1.71	3.27	6.19
Dataset	% sampled	Egonet Sample Size						
		15	**31**	**62**	**125**	**250**	**500**	**1K**
FB: UVA	% nodes	7.42	14.50	25.66	41.60	60.11	75.70	86.58
	% edges	2.20	4.41	8.47	15.97	28.28	44.70	63.90
FB: UCSD	% nodes	5.73	11.12	20.30	33.69	51.16	67.93	81.37
	% edges	1.92	3.93	7.77	14.21	25.91	42.10	61.34

Table 2: Uniform sampling without replacement. Total % Nodes Sampled in the graph when including all egos and neighbors. Total % Edges Sampled in the graph when including all edges between egos and neighbors.

formance, even for very large graphs.

5.3.2 *Effect of sampling method*

We now consider the effect of the sampling method on the performance of our estimators. In this evaluation, the subgraphs of interest will be all 5-node undirected subgraphs in Fig. 6c. For reasons of spatial complexity we will only use the Role Occupancy estimator. To give an example of the challenges involved with the Unique Counting estimator, there are 627 and 10 trillion 5-node subgraphs of type 1 and type 2 correspondingly in the graph "web-Google." The UC estimator requires that we keep track of every distinct subgraph instance which is impractical.

In Fig. 7 we compare four sampling methods: (1) Uniform Independence Sampling (UIS), (2) Random Walk (RW), (3) Thinned Random Walk, where we collect one egonet every 30 samples in the random walk, and (4) Weighted Independent Sampling (WIS) with the weight of each node set equal to its degree. The NRMSE is shown for the estimation of the subgraph count for all 5-node undirected subgraphs of Fig. 6c, in graphs "soc-Slashdot," "email-EuAll," and "web-Google."

We observe that for UIS, subgraphs with higher number of nodes in observable roles have lower estimation error. RW samples, however, yield estimation error that is sometimes more than one order of magnitude smaller than UIS in the network soc-Slashdot. The intuition behind this observation is that RW samples are biased toward higher degree nodes which contain a proportionally larger number of subgraph counts. Since we appropriately reweight for the bias, we get a much better (unbiased) estimate of subgraph counts compared to UIS. We should note that the performance boost for RW samples depends on the structure of the network. In cases such as web-Google and email-EuAll, where consecutive RW egonet samples are very correlated and mixing in the network is slow, the performance of RW is comparable with that of UIS. In the latter cases, we observe that by applying thinning in the random walk (every 30 samples) we

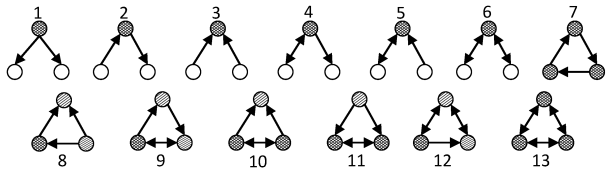

(a) 3-node directed subgraphs

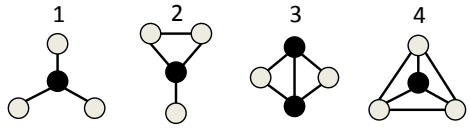

(b) 4-node undirected subgraphs

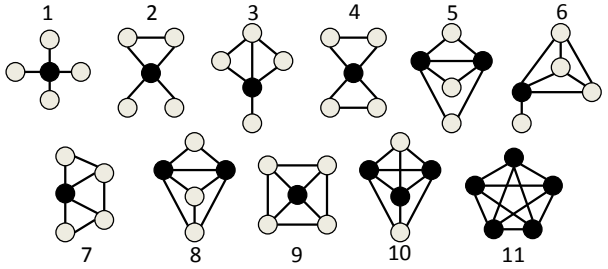

(c) 5-node undirected subgraphs

Figure 6: Subgraphs used during simulations. For 4 and 5-node subgraphs, the nodes in black correspond to observable orbits. 3-node subgraphs have multiple observable semi-spanning orbits and the corresponding nodes are colored with a different pattern.

reduce the correlation of consecutive node to such a degree that we can reach the NRMSE of a WIS sample.

5.3.3 *Comparison with PSRW*

To assess relative performance, we compare our Role Occupancy estimator with the state-of-the-art method PSRW [53], which was shown to significantly outperform FANMOD and GUISE. We received from the authors of [53] source code that implements SRW and PSRW for all 3-node directed subgraphs and 4-node undirected subgraphs.[7]

Fig. 8 shows the error of the two estimators when estimating all 4-node undirected subgraphs of Fig. 6b for networks "soc-Slashdot" and "Amazon." In both networks, RO outperforms PSRW by at least an order of magnitude for subgraph IDs 2, 3, and 4. On the other hand for subgraph ID 1, RO underperforms until sample size $n = 10,000$. We also observe that whereas in "soc-Slashot" network thinning does not yield any performance improvements, in the "Amazon" network thinning the Random Walk every 30 samples yields a significant improvement.

Fig. 9 shows the error of the two estimators when estimating all [8] 3-node undirected subgraphs (see Fig. 6a) for the network "youtube-links." We observe that the RO method always outperforms PSRW by a significant margin. RO requires approximately 500-1000 egonet samples to reach NRMSE at or below 10^{-1} for all subgraph IDs. On the other hand, in cases of subgraph IDs 7, 8, and 9, the NRMSE for

[7]The authors of [53] claim that their method becomes impractically slow for subgraphs of larger size.

[8]NRMSE for subgraph type 8 is not shown in Fig. 9 due to similarity with type 9 and economy of space.

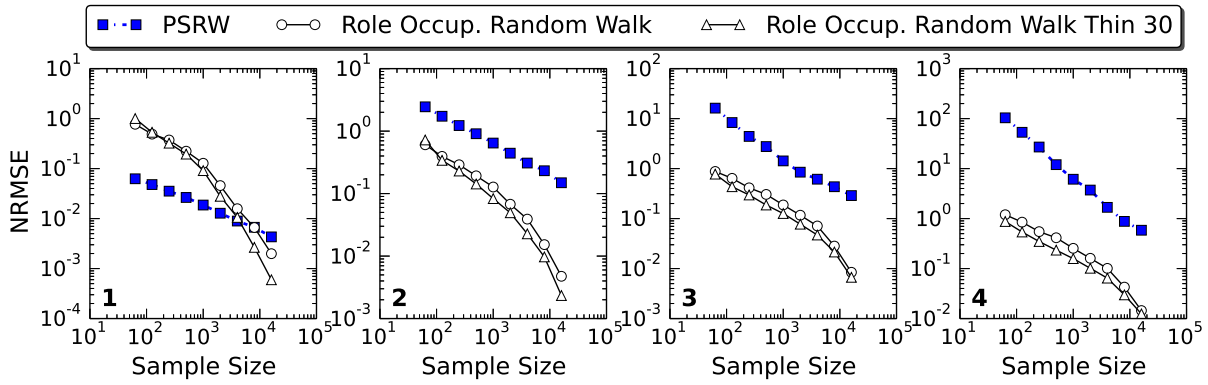

(a) soc-Slashdot

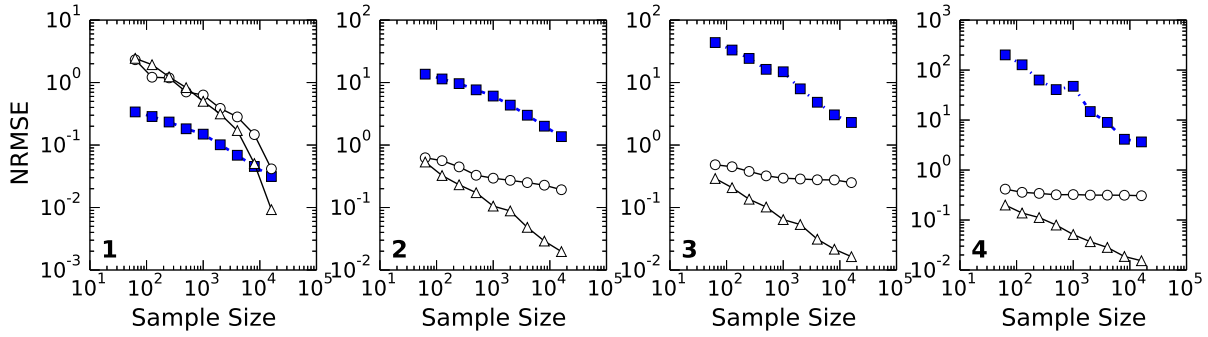

(b) amazon

Figure 8: NRMSE of PSRW and Role Occupancy for the estimation of all 4-node undirected subgraphs supported by our egocentric method

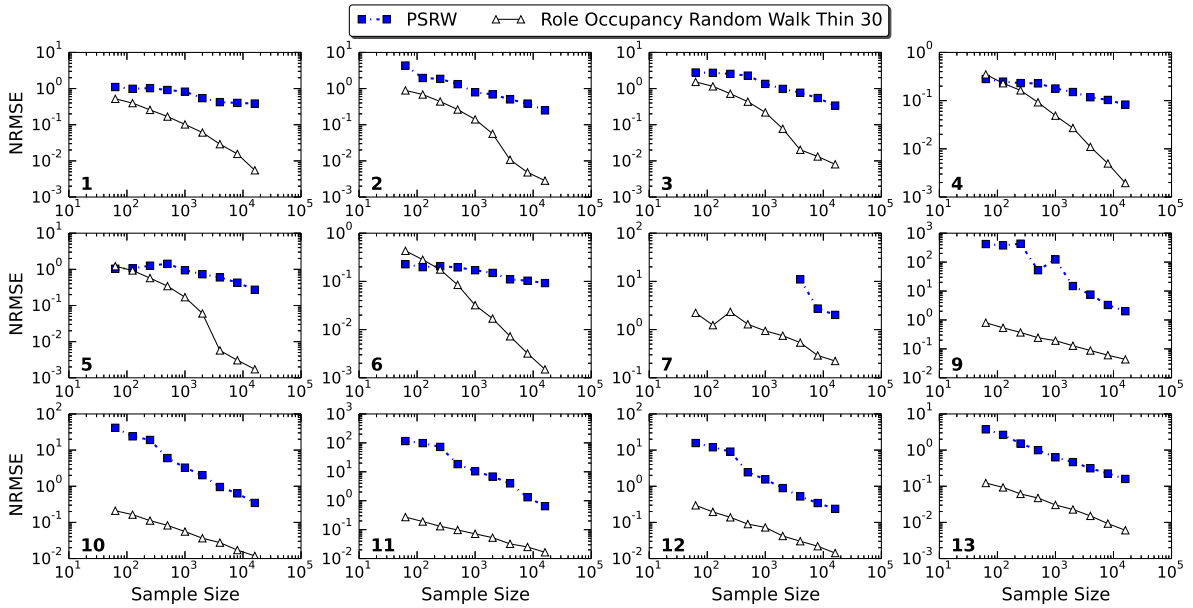

Figure 9: NRMSE of PSRW and Role Occupancy (our method) for the estimation of all 3-node directed subgraphs in the network "youtube-links"

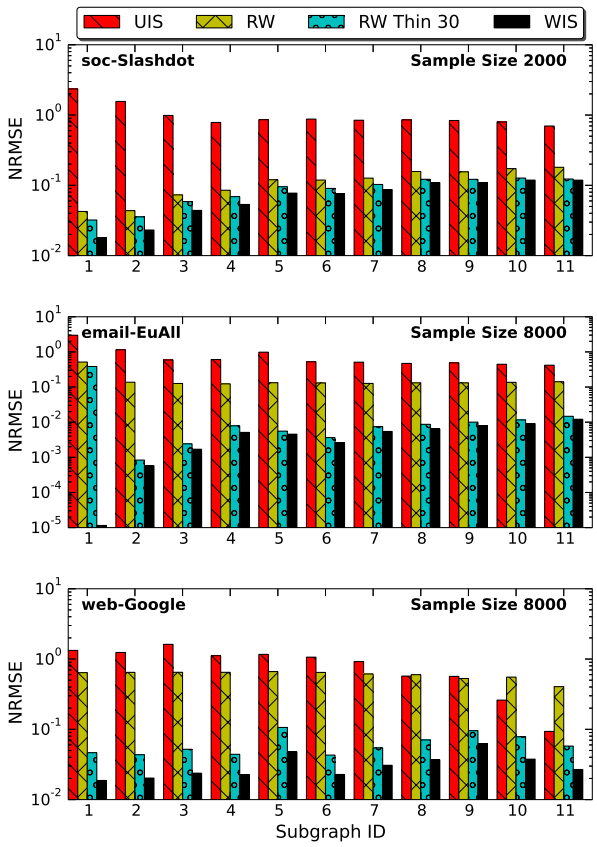

Figure 7: NRMSE for the estimation of all 5-node undirected subgraph IDs (see Fig. 6c) in three different graphs averaged over 100 simulation iterations. For each subgraph ID, we list the NRMSE for sampling methods UIS, RW, RW Thin 30, WIS.

PSRW is as high as $5-10$ after $16K$ samples.

6. APPLICATION TO FACEBOOK

6.1 Dataset Description

In previous work [15], we collected a representative sample of ≈ 1 million unique Facebook (FB) users by crawling the social graph using a Metropolis Hasting Random Walk (MHRW) method. Subsequently we collected the egonets for $36,628$ unique nodes that were randomly selected from the MHRW sample. This sub-sampling eliminates the correlation of consecutive nodes in the same crawl, similarly to the "Thin 30" sample in Section 5.3.2. The representativeness of this data has been validated against true random samples from the Facebook taken during the same period [15,16]. This sample closely approximates a uniform, with replacement sample of egonets from the publicly visible FB graph. In this sample all neighborhoods are uniquely labeled which allows for estimation using either the role occupancy or unique counting estimators. Due to reasons of space complexity, here we use the role occupancy estimators. We use the population size $N = 240M$ which was estimated for this dataset by [31] and agrees with the statistics reported by Facebook during the collection of the dataset (April 2009).

We complement this egonet sample with gender attributes

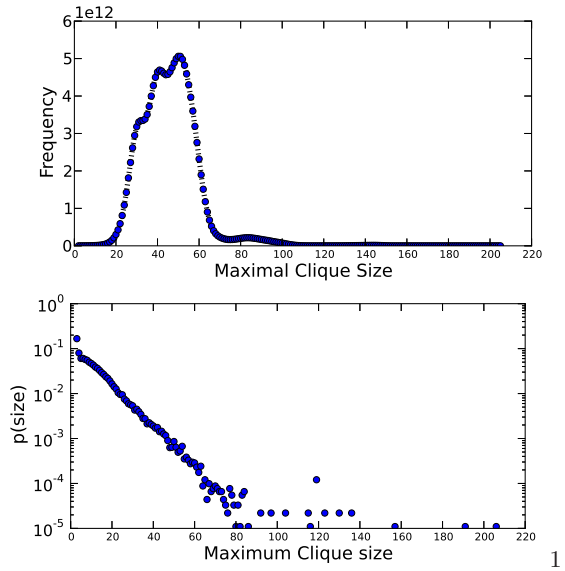

Figure 10: Estimated clique size distribution (Facebook social graph '09); top panel shows Role Occupancy estimates of maximal clique frequency, bottom panel shows the distribution of maximum clique sizes by ego.

for each user. We were able to fetch the publicly declared user-declared gender for 90% of sampled users by crawling the url at *http://graph.facebook.com/userid*. Additionally, we classified another 9.5% by a majority rule that uses the first name of each user and a database of the number of times that first names were assigned to males and females. We first used the list of first names from the US Social Security records. If there was no match we then used the list of first names from the population of Facebook users with declared gender. Last, we used [2] to predict the gender for the remaining 0.5% users with a Naive Bayes classifier, based on the letter composition of first names.

6.2 Results

The top panel of Fig. 10 shows the estimated distribution of maximal clique sizes over the entire FB social graph. The FB graph is known to be highly clustered, and it indeed contains many large cliques: the modal clique size is 50, with the largest observed clique containing over 205 individuals. Interestingly, the form of the clique distribution is neither monotone nor unimodal; significant peaks occur at 32, 41, and 50, with a minor mode near 84. This suggests substantial heterogeneity in the mechanisms of clique formation, a point underscored by our findings regarding gender (see below).

Rather more order is shown in the distribution of maximum clique sizes by ego (i.e., the largest clique to which a given individual belongs). (Fig. 10, bottom panel.) This shows a monotone distribution with a stable exponential decay over the range that is well-supported by our data. Membership in moderate to large cliques is thus quite rare, despite their prevalence in the FB graph.

Beyond size distributions, our estimators allow us to examine how the composition of cliques varies across the FB graph. Fig. 11 shows the estimated gender composition of FB cliques for all cliques (panel a) and cliques of varying order (panels b-h). The X axis in each panel indicates the

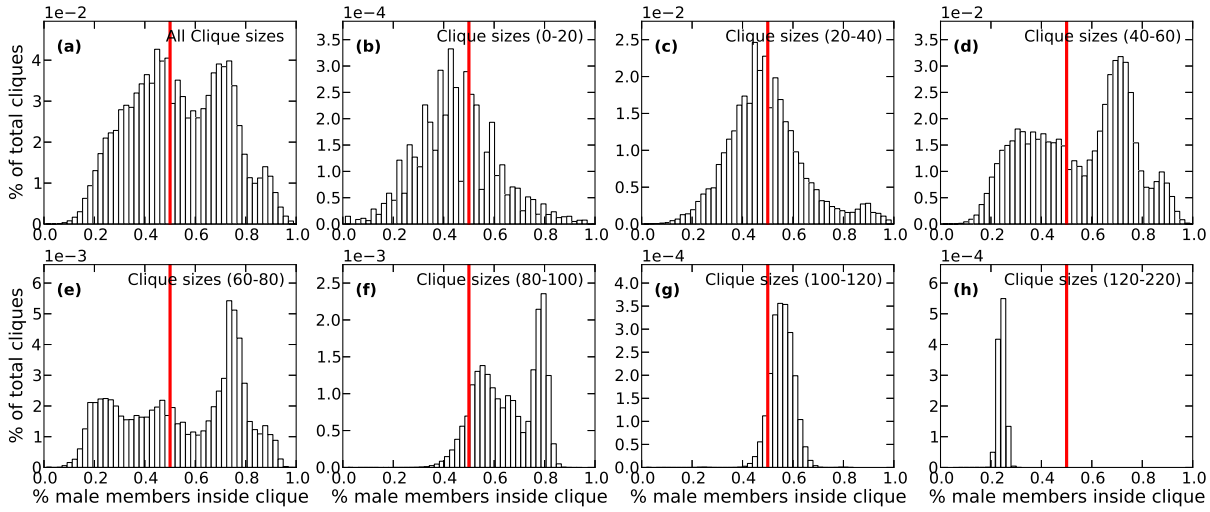
(a) All Clique sizes
(b) Clique sizes (0-20)
(c) Clique sizes (20-40)
(d) Clique sizes (40-60)
(e) Clique sizes (60-80)
(f) Clique sizes (80-100)
(g) Clique sizes (100-120)
(h) Clique sizes (120-220)
% of total cliques
% male members inside clique

Figure 11: Role Occupancy estimates for gender composition of maximal cliques, by order. (FB social graph '09)

fraction of clique members who are male, from 0 (entirely female) to 1 (entirely male); a vertical reference line indicates gender parity. Our results provide clear evidence for strong heterogeneity in the makeup of FB cliques. We see several distinct modes with characteristic gender frequencies, that occur over specific size ranges. These include: a "small equal clique" mode of near-parity cliques of size 0–40; a 70–80% male mode for cliques of size 40-100; a 60-80% female mode for cliques of order 40–80; a second near-parity mode over the small range of sizes 100–120; and a strongly female dominated mode of very large cliques (sizes > 120). Although our data does not allow us to establish the mechanisms underlying these modes, we speculate that each is the result of a particular collection of social settings (e.g., fraternities or sororities, family groups, schools, or work organizations) that acts as a focus [13] for tie formation. Systematic variation in the gender composition of these settings then leads to corresponding variation in clique composition. Interestingly, our findings do not corroborate the predictions of [38] regarding the relationship of clique size to gender homogeneity based on their analysis of face-to-face groups: while they posit a strongly negative relationship between heterogeneity and group size, we find that the FB graph supports a large fraction of near-parity cliques at even quite large orders. While it is also true that extremely gender-homogeneous cliques become relatively more prevalent at large orders (versus small ones), the phenomenon appears to be uneven and size-specific. Neither do we observe the power-law decay in group sizes reported by [38] for naturally occurring groups. Since these prior results were based on observations of cliques in public, face-to-face settings, this lack of replication does not necessarily call into question the validity of the authors' conclusions in their original context; however, it does underscore that the formation of friendship cliques in OSNs may operate very differently from the sorts of groups examined in previous studies.

6.3 Dataset Release

We make available at [1] an anonymized sample of the Facebook egonets that were used in the analysis of this Section. In particular, the dataset contains 36,628 egonets, the gender for all egos and neighbors (a total of $\approx 5,6M$ users), and the maximal clique computation without attributes and by gender for all egonets.

7. CONCLUSION

In this paper, we introduced a family of novel unbiased estimators of subgraph counts based on egocentric network samples. We presented two techniques, one of which exploits labeling of nodes (UC) and one which does not require this information (RO). Both techniques are parallelizable, and suitable for use with large graphs. We evaluated estimator performance via simulated sampling from real-world graphs, showing that both proposed techniques work well and that UC generally outperforms RO as the sample "saturates" the graph. We showed that our techniques match or surpass the state-of-the-art method PSRW for subgraph counting and that they can be used for much larger subgraphs. Finally, we demonstrated an application of our estimators to clique composition in OSNs. We applied our methodology to egocentric samples collected in Facebook, which we complemented with gender information, allowing us to estimate the joint size and composition distribution of FB cliques with respect to gender. Our results underscore important differences between online and (previously reported) offline group structure, and provide evidence for strong gender heterogeneity in the makeup of FB cliques.

8. REFERENCES

[1] Facebook Egonets Sample. `http://odysseas.calit2.uci.edu/doku.php/public:osn_datasets`

[2] Gender Predictor. `https://github.com/sholiday/genderPredictor`.

[3] SNAP Graph Library. `http://snap.stanford.edu/data/`.

[4] A.-L. Barabasi and Z. N. Oltvai. Network biology: understanding the cell's functional organization. *Nature reviews genetics*, 5(2):101–113, 2004.

[5] H. R. Bernard, P. D. Killworth, and L. Sailer. Informant accuracy in social network data iv: A comparison of clique-level structure in behavioral and

cognitive network data. *Social Networks*, 2(3):191–218, 1980.

[6] M. Bhuiyan, M. Rahman, M. Al Hasan, et al. Guise: Uniform sampling of graphlets for large graph analysis. In *Data Mining (ICDM), 2012 IEEE 12th International Conference on*, pages 91–100. IEEE, 2012.

[7] K. R. W. Brewer and M. Hanif. *Sampling with Unequal Probabilities*. Springer-Verlag, 1983.

[8] R. S. Burt. Network items and the general social survey. *Social networks*, 6(4):293–339, 1984.

[9] S. M. Caldeira, T. P. Lobao, R. F. S. Andrade, A. Neme, and J. V. Miranda. The network of concepts in written texts. *The European Physical Journal B-Condensed Matter and Complex Systems*, 49(4):523–529, 2006.

[10] P. J. Carrington, J. Scott, and S. Wasserman. *Models and methods in social network analysis*. Cambridge university press, 2005.

[11] J. Chen, W. Hsu, M. L. Lee, and S.-K. Ng. Nemofinder: Dissecting genome-wide protein-protein interactions with meso-scale network motifs. In *Proceedings of the 12th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 106–115. ACM, 2006.

[12] D. Eppstein and D. Strash. Listing all maximal cliques in large sparse real-world graphs. In *Experimental Algorithms*, pages 364–375. Springer, 2011.

[13] S. Feld. The focused organization of social ties. *American Journal of Sociology*, 1986:1015–1035, 1981.

[14] O. Frank and D. Strauss. Markov graphs. *Journal of the American Statistical Association*, 81:832–842, 1986.

[15] M. Gjoka, M. Kurant, C. T. Butts, and A. Markopoulou. Walking in Facebook: A Case Study of Unbiased Sampling of OSNs. In *Proceedings of IEEE INFOCOM '10*.

[16] M. Gjoka, M. Kurant, C. T. Butts, and A. Markopoulou. Practical recommendations on crawling online social networks. *IEEE JSAC on Measurement of Internet Topologies*, 29(9):1872–1892, 2011.

[17] M. Gjoka, E. Smith, and C. T. Butts. Estimating Clique Composition and Size Distributions from Sampled Network Data. In *Proceedings of IEEE NetSciCom '14*, Toronto, Canada, May 2014.

[18] R. Gould and R. Fernandez. Structures of mediation: A formal approach to brokerage in transaction networks. *Sociological Methodology*, 19:89–126, 1989.

[19] H. M. Grindley, P. J. Artymiuk, D. W. Rice, and P. Willett. Identification of tertiary structure resemblance in proteins using a maximal common subgraph isomorphism algorithm. *Journal of molecular biology*, 1993.

[20] J. A. Grochow and M. Kellis. Network motif discovery using subgraph enumeration and symmetry-breaking. In *Research in Computational Molecular Biology*, pages 92–106. Springer, 2007.

[21] M. H. Hansen and W. N. Hurwitz. On the theory of sampling from finite populations. *The Annals of Mathematical Statistics*, 14(4), 1943.

[22] S. J. Hardiman, P. Richmond, and S. Hutzler. Calculating statistics of complex networks through random walks with an application to the on-line social network Bebo. *The European Physical Journal B*, 2009.

[23] T. Hočevar and J. Demšar. A combinatorial approach to graphlet counting. *Bioinformatics*, 30(4):559–565, 2014.

[24] P. W. Holland and S. Leinhardt. Transitivity in structural models of small groups. *Comparative Group Studies*, 2:107–124, 1971.

[25] P. W. Holland and S. Leinhardt. The statistical analysis of local structure in social networks. *Sociological Methodology*, 6:1–45, 1975.

[26] P. W. Holland and S. Leinhardt. An exponential family of probability distributions for directed graphs (with discussion). *Journal of the American Statistical Association*, 76(373):33–50, 1981.

[27] D. R. Hunter, P. N. Krivitsky, and M. Schweinberger. Computational statistical methods for social network analysis. *Journal of Computational and Graphical Statistics*, 21:856–882, 2012.

[28] I. Jansson. Clique structure in school class data. *Social Networks*, 19(3):285–301, 1997.

[29] Z. R. Kashani, H. Ahrabian, E. Elahi, A. Nowzari-Dalini, E. S. Ansari, S. Asadi, S. Mohammadi, F. Schreiber, and A. Masoudi-Nejad. Kavosh: a new algorithm for finding network motifs. *BMC bioinformatics*, 10(1):318, 2009.

[30] N. Kashtan, S. Itzkovitz, R. Milo, and U. Alon. Efficient sampling algorithm for estimating subgraph concentrations and detecting network motifs. *Bioinformatics*, 20(11):1746–1758, 2004.

[31] L. Katzir, E. Liberty, and O. Somekh. Estimating Sizes of Social Networks via Biased Sampling. In *WWW*, 2011.

[32] E. D. Kolaczyk. *Statistical Analysis of Network Data*. 2009.

[33] D. Krackhardt and M. Kilduff. Structure, culture and simmelian ties in entrepreneurial firms. *Social networks*, 24(3):279–290, 2002.

[34] M. Kurant, M. Gjoka, C. T. Butts, and A. Markopoulou. Walking on a Graph with a Magnifying Glass: Stratified Sampling via Weighted Random Walks. In *Proceedings of ACM SIGMETRICS '11*, San Jose, CA, June 2011.

[35] L. Lovász. Random walks on graphs: A survey. *Combinatorics, Paul Erdos is Eighty*, 2(1):1–46, 1993.

[36] P. V. Marsden. Network data and measurement. *Annual review of sociology*, pages 435–463, 1990.

[37] P. V. Marsden. Recent developments in network measurement. *Models and methods in social network analysis*, 8:30, 2005.

[38] B. H. Mayhew, J. M. McPherson, M. Rotolo, and L. Smith-Lovin. Sex and race homogeneity in naturally occurring groups. *Social Forces*, 1995.

[39] R. Milo, S. Shen-Orr, S. Itzkovitz, N. Kashtan, D. Chklovskii, and U. Alon. Network motifs: simple building blocks of complex networks. *Science*, 298(5594):824–827, 2002.

[40] A. Mislove, M. Marcon, K. Gummadi, P. Druschel,

and B. Bhattacharjee. Measurement and analysis of online social networks. In *IMC '07*.

[41] S. Omidi, F. Schreiber, and A. Masoudi-Nejad. Moda: an efficient algorithm for network motif discovery in biological networks. *Genes & genetic systems*, 84(5):385–395, 2009.

[42] N. Pržulj. Biological network comparison using graphlet degree distribution. *Bioinformatics*, 23(2):e177–e183, 2007.

[43] F. Schreiber and H. Schwöbbermeyer. Mavisto: a tool for the exploration of network motifs. *Bioinformatics*, 21(17):3572–3574, 2005.

[44] T. A. B. Snijders, P. E. Pattison, G. L. Robins, and M. S. Handcock. New specifications for exponential random graph models. *Sociological Methodology*, 36, 2006.

[45] E. S. Spiro, R. M. Acton, and C. T. Butts. Extended structures of mediation: Re-examining brokerage in dynamic networks. *Social Networks*, 35:130–143, 2013.

[46] W. Stein et al. Sage: Open source mathematical software. *7 December 2009*, 2008.

[47] F. Stentiford and A. Bamidele. Image recognition using maximal cliques of interest points. In *Proc. of IEEE ICIP '10*.

[48] D. M. Strickland, E. Barnes, and J. S. Sokol. Optimal protein structure alignment using maximum cliques. *Operations research*, 53(3):389–402, 2005.

[49] S. K. Thompson. *Sampling*. Wiley, New York, second edition, 2002.

[50] Y. Tille and A. Matei. *sampling: Survey Sampling*, 2012. R package v2.5.

[51] A. Traud, P. Mucha, and M. Porter. Social Structure of Facebook Networks. *Arxiv preprint arXiv:1102.2166*, 2011.

[52] B. Viswanath, A. Mislove, M. Cha, and K. Gummadi. On the evolution of user interaction in facebook. In *Proc. of WOSN '09*.

[53] P. Wang, J. Lui, B. Ribeiro, D. Towsley, J. Zhao, and X. Guan. Efficiently estimating motif statistics of large networks. *ACM Transactions on Knowledge Discovery from Data (TKDD)*, 9(2):8, 2014.

[54] S. Wasserman and K. Faust. *Social Network Analysis: Methods and Applications*. Cambridge University Press, Cambridge, 1994.

[55] S. Wasserman and K. Faust. *Social network analysis: Methods and applications*. Cambridge university press, 1994.

[56] S. Wernicke and F. Rasche. Fanmod: a tool for fast network motif detection. *Bioinformatics*, 22(9):1152–1153, 2006.

[57] E. Wong, B. Baur, S. Quader, and C.-H. Huang. Biological network motif detection: principles and practice. *Briefings in bioinformatics*, 13(2):202–215, 2012.

[58] O. N. Yaveroğlu, S. M. Fitzhugh, M. Kurant, A. Markopoulou, C. T. Butts, and N. Pržulj. `ergm.graphlets`: A package for ERG modeling based on graphlet statistics. *Journal of Statistical Software*, 65(12), 2015.